

CLAIMS

What is claimed is:

1. A method of recovery from a data storage system failure in a data storage system having a host computer writing data to a first storage unit associated with a first storage controller synchronously mirroring the data to a second storage unit associated with a second storage controller asynchronously mirroring the data to a third storage unit, the method comprising:

detecting a failure associated with the first storage unit;

writing data updates directly to the second storage unit;

correcting the failure associated with the first storage unit; and

synchronously mirroring the data updates from the second storage unit to the first storage unit.

2. The method of claim 1 further comprising suspending the synchronous mirroring of the data from the first storage unit to the second storage unit upon detection of the failure associated with the first storage unit.

3. The method of claim 1 further comprising suspending the asynchronous mirroring of the data updates from the second storage unit to the third storage unit while synchronously mirroring the data updates from the second storage unit to the first storage unit.

4. The method of claim 3 further comprising:

re-establishing the synchronous mirroring of the data from the first storage unit to the second storage unit;

resuming the writing of the data from the host computer to the first storage unit; and

re-establishing the asynchronous mirroring of data from the second storage unit to the third storage unit.

5. The method of claim 1 wherein the asynchronous mirroring of the data comprises:

creating a first map associated with the second storage unit identifying tracks storing mirrored data which has been received on the second storage unit by the synchronous mirroring of the data from the first storage unit; and

continuously transmitting batches of the mirrored data on tracks identified by the first map from the second storage unit to the third storage unit.

6. The method of claim 5 wherein the synchronous mirroring of the data updates from the second storage unit to the first storage unit comprises:

creating a second map associated with the second storage unit identifying the tracks storing changed data from a point in time where the host computer began writing directly to the second storage unit;

copying to the first storage unit the changed data on the tracks identified by the second map; and

synchronously mirroring to the first storage unit further changed data written to the second storage unit by the host computer until a full duplex state between the first storage unit and the second storage unit is reached.

7. The method of claim 6 wherein the copying of the changed data on the tracks identified by the second map to the first storage unit comprises:

preserving a copy of the second map at a point in time at which copying the changed data to the first storage unit begins;

re-setting the second map to a clear state and thereafter identifying the tracks containing the changed data in the second map;

comparing the first map to the second map and adding the identity of the tracks containing the data updates, as identified by the first map but not identified by the second map, to the second map;

comparing the preserved copy of the second map to the first map and adding the identity of the tracks containing the changed data, as identified by the preserved copy of the second map but not identified by the first map, to the first map; and

copying the data updates on the tracks then identified by the first map from the second storage unit to the first storage unit.

8. The method of claim 6 wherein the synchronous mirroring to the first storage unit of the further updates written to the second storage unit by the host computer comprises:

continuing to identify the tracks having the data updates with the first map and with the second map; and

synchronously copying the data updates on the tracks identified by the first map from the second storage unit to the first storage unit.

9. The method of claim 6 further comprising the following steps when the full duplex state between the first storage unit and the second storage unit is reached:

terminating the use of the first map to identify the tracks associated with the second storage unit storing the data updates; and

continuing to use the second map to identify the tracks associated with the second storage unit storing changed data.

10. The method of claim 6 wherein the writing of the data by the host computer is quiesced after the full duplex state between the first storage unit and the second storage unit is reached.

11. The method of claim 6 wherein re-establishing the asynchronous mirroring of the data from the second storage unit to the third storage unit comprises:

comparing the second map to the first map and adding the identity of the tracks containing the changed data, as identified by the second map but not identified by the first map, to the first map; and

asynchronously mirroring the data updates on the tracks then identified by the first map from the second storage unit to the third storage unit.

12. A system for copying stored data and having the ability to recover from a failure, comprising:

a second storage controller associated with a second storage unit, the second storage controller synchronously receiving mirrored data from a first storage controller associated with a host computer and a first storage unit, the second storage controller further asynchronously mirroring the data to a third storage controller associated with a third storage unit;

means for detecting a failure associated with the first storage unit;

means for writing data updates from the host computer to the second storage controller upon detection of the failure associated with the first storage unit;

means for synchronously mirroring the data updates from the second storage unit to the first storage unit upon correction of the failure associated with the first storage unit.

13. The system for copying stored data of claim 12 wherein the second storage controller further comprises a first map identifying tracks storing the mirrored data which has been received on the second storage unit by the synchronous mirroring of the data from the first storage unit.

14. The system for copying stored data of claim 12 wherein the second storage controller further comprises a second map identifying the tracks storing changed data from a point in time where the host computer begins writing the data updates directly to the second storage unit.

15. The system for copying stored data of claim 12 wherein the second storage controller further comprises:

a copy of the second map made at a point in time the synchronous mirroring of the data updates to the first storage unit begins;

the second map set to a clear state at the point in time the synchronous mirroring of the data updates to the first storage unit begins and thereafter identifying the tracks containing new changed data.

16. The system for copying stored data of claim 15 wherein the second storage controller further comprises means for comparing the tracks identified by the first map to the tracks identified by one of the second map, and the preserved copy of the second map.

17. The system for copying stored data of claim 16 wherein the means for synchronously mirroring the data updates from the second storage unit to the first storage unit causes the following steps to be taken:

comparing the first map to the re-set second map and adding the identity of the tracks containing the data updates, as identified by the first map but not identified by the second map, to the second map;

comparing the preserved copy of the second map to the first map and adding the identity of the tracks containing the changed data, as identified by the preserved copy of the second map but not identified by the first map, to the first map; and

copying the data updates on the tracks then identified by the first map from the second storage unit to the first storage unit.

18. The system for copying stored data of claim 12 wherein the second storage controller further comprises means for terminating the asynchronous mirroring of the data to the third storage controller associated with the third storage unit upon the commencement of the synchronous mirroring of the data updates from the second storage unit to the first storage unit.

19. The system for copying stored data of claim 18 wherein the second storage controller further comprises:

means for re-establishing the asynchronous mirroring of the data to the third storage controller associated with the third storage unit; and

means for synchronizing the data stored on the third storage unit with the data stored on the second storage unit.

20. An article of manufacture for use in programming a data storage system to recover from a failure, the data storage system having a host computer writing data to a first storage unit associated with a first storage controller synchronously mirroring the data to a second storage unit associated with a second storage controller asynchronously mirroring the data to a third storage unit, the article of manufacture comprising a storage medium having logic embedded therein to cause components of the data storage system to:

detect a failure associated with the first storage unit;

write data updates directly to the second storage unit;

correct the failure associated with the first storage unit; and

synchronously mirror the data updates from the second storage unit to the first storage unit.

21. The article of manufacture of claim 20 wherein the logic further causes components of the data storage system to suspend the synchronous mirroring of the data from the first storage unit to the second storage unit upon detection of the failure associated with the first storage unit.
 22. The article of manufacture of claim 20 wherein the logic further causes components of the data storage system to suspend the asynchronous mirroring of the data updates from the second storage unit to the third storage unit while synchronously mirroring the data updates from the second storage unit to the first storage unit.
23. The article of manufacture of claim 22 wherein the logic further causes components of the data storage system to:
- re-establish the synchronous mirroring of the data from the first storage unit to the second storage unit;
 - resume the writing of the data from the host computer to the first storage unit; and
 - re-establish the asynchronous mirroring of data from the second storage unit to the third storage unit.

24. The article of manufacture of claim 20 wherein the logic further causes asynchronous mirroring of the data from the second storage unit to the third storage unit by:

creating a first map associated with the second storage unit identifying tracks storing mirrored data which has been received on the second storage unit by the synchronous mirroring of the data from the first storage unit; and

periodically transmitting batches of the mirrored data on tracks identified by the first map from the second storage unit to the third storage unit.

25. The article of manufacture of claim 24 wherein the logic further causes synchronous mirroring of the data updates from the second storage unit to the first storage unit by:

creating a second map associated with the second storage unit identifying the tracks storing changed data from a point in time where the host computer began writing directly to the second storage unit;

copying to the first storage unit the changed data on the tracks identified by the second map; and

synchronously mirroring to the first storage unit further changed data written to the second storage unit by the host computer until a full duplex state between the first storage unit and the second storage unit is reached.

26. The article of manufacture of claim 25 wherein the logic further causes copying of the changed data on the tracks identified by the second map to the first storage unit by:

preserving a copy of the second map at a point in time at which copying the changed data to the first storage unit begins;

re-setting the second map to a clear state and thereafter identifying the tracks containing the changed data in the second map;

comparing the first map to the second map and adding the identity of the tracks containing the data updates, as identified by the first map but not identified by the second map, to the second map;

comparing the preserved copy of the second map to the first map and adding the identity of the tracks containing the changed data, as identified by the preserved copy of the second map but not identified by the first map, to the first map; and

copying the data updates on the tracks then identified by the first map from the second storage unit to the first storage unit.

27. The article of manufacture of claim 25 wherein the logic further causes synchronous mirroring to the first storage unit further updates written to the second storage unit by the host computer by:

continuing to identify tracks having the data updates with the first map and with the second map; and

synchronously copying the data updates on the tracks identified by the first map from the second storage unit to the first storage unit.

28. The article of manufacture of claim 25 wherein the logic further causes components of the data storage system to take the following steps when the full duplex state between the first storage unit and the second storage unit is reached:

terminate the use of the first map to identify the tracks associated with the second storage unit storing the data updates; and

continue to use the second map to identify the tracks associated with the second storage unit storing changed data.

29. The article of manufacture of claim 25 wherein the logic further causes components of the data storage system to stop the writing of data by the host computer after the full duplex state between the first storage unit and the second storage unit is reached.

30. The article of manufacture of claim 25 wherein the logic further causes the re-establishing of the asynchronous mirroring of the data from the second storage unit to the third storage unit by:

comparing the second map to the first map and adding the identity of the tracks containing the changed data, as identified by the second map but not identified by the first map, to the first map; and

asynchronously mirroring the data updates on the tracks then identified by the first map from the second storage unit to the third storage unit.